Technical Progress Report 11/1/95 – 1/31/96
Construction of a Connectionist Network Supercomputer
University of California, Berkeley
ONR URI Grant No. N00014-92-J-1617

1 Abstract

This report presents a summary of the technical status for the period 11/1/95–1/31/96.

In the previous report, we reported the running of a complete neural network training on our SPERT neurocomputing board. This quarter we moved to the next phase of the project—distribution of our neurocomputing boards to other researchers. We reached a major milestone with the release of our neurocomputing board and software to another research laboratory. In addition, we continued to make good progress in high level software, and analog auditory preprocessors for speech recognition.

2 Technical Status

2.1 The SPERT System

In this quarter we made a concentrated effort to prepare the release of SPERT neurocomputing board to other sites. On the hardware side, we developed rigorous board testing and verification procedures. We also made several minor modifications to the boards to aid its reliability. On the software side, we performed a major cleanup, test, and documentation of the SPERT development environment and runtime system. In January, only a couple of weeks behind schedule, we installed a SPERT system in the laboratory of Hervé Bourland, a speech researcher at Faculté Polytechnique de Mons, Mons, Belgium.

We currently have eight complete running boards at Berkeley and 15 more nearly completed. A few of the boards are reserved for "production" experimental speech recognition work; these are in constant use by our speech researchers and have already had a significant effect on the quality of the research. In the next quarter we plan to distribute approximately 15 boards to other sites. We have also begun procurement for an additional 32 boards.

This quarter we presented a paper on SPERT to the 1995 Neural Information Processing Systems Conference (NIPS 95). This, the most prestigious of the neural networks conferences, accepts very few implementation papers and we felt fortunate to have two accepted (one on SPERT and another on our work on Analog VLSI pre-processors). We also prepared a manuscript to appear as a feature article in IEEE Computer this spring.

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • RIVERSIDE • SAN DIEGO • SAN FRANCISCO

SANTA BARBARA • SANTA CRUZ

COLLEGE OF ENGINEERING
DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCES
COMPUTER SCIENCE DIVISION

BERKELEY, CALIFORNIA 94720

February 9, 1996

Dr. Nick Bottka
Electronics Division, ONR 312
Office of Naval Research
Scientific Office Code 1114SE
800 North Quincy Street
Arlington, VA 22217-5660

Reference: ONR URI Grant No.: 144-92-J-1617

Dear Dr. Bottka:

Enclosed you will find the quarterly progress report for the research project for Professor John Wawrzynek: "Construction of a Connectionist Network Supercomputer," ONR URI Grant No.: N00014-92-J-1617. This report covers the research done during the performance period 11/1/95-1/31/96. Please direct any mail or questions regarding this contract to me.

Thank you for your continued support of this project.

Sincerely,

Theresa Lessard-Smith
Project Coordinator

cc: Administrative Grants Officer, Seattle
Director, Naval Research Laboratory
Defense Technical Information Center
Sponsored Projects Office - UC Berkeley
file

2.2 High-Level Software

There was also significant progress in the area of high level programming for parallelism. Sather release 1.0.8 has been well received and is being used for an advanced parallel programming course on campus. The Myrinet system for parallel Sather has been installed and is in routine daily use. Our group was awarded a grant from Lawrence Livermore National Lab supporting use of their Meiko CS2 super computer. In this quarter we successfully ported pSather to the Meiko and this provides for much greater peak capacity. Good progress was made on the automatic mapping facilities for the ICSIM connectionist simulator.

2.3 Analog VLSI pre-processors.

As described in our last progress report, active research on our speaker-independent, telephone-quality isolated-word speech recognition system is drawing to a close; this system combines our analog VLSI sound pre-processing system with a Hidden-Markov-Model-based speech recognizer that uses Multi-Layer Perceptrons for phoneme classification.

This quarter, we focused on writing a full-length research publication describing the system, and submitted it to a Special Issue on Neuromorphic Systems of the journal Analog Integrated Circuits and Signal Processing. The article is currently in review; a copy of the article is included with our status report, and is also available on the World Wide Web at

`http://www.pcmp.caltech.edu/anaprose/lazzaro/recog.ps.Z`

We also intend to present the work described in the article at "Machines that Learn" conference in April of this year in Utah.

The main points of our journal article were:

1. Speech recognition performance of systems using our analog VLSI pre-processors is adequate for use in applications (4.1% error (speaker-independent) on a 13-word task), but is currently inferior to state-of-the-art traditional front-ends (1.8% error on the same task).
2. In the short run, we believe these speech pre-processing techniques may find application in micropower speech recognition systems, where our circuit techniques have an important advantage over competitive approaches: real-time performance with microwatt power consumption.
3. In the long run, we believe biologically-based front-end processing has the potential to offer superior recognition performance to conventional front-ends. We detail promising research directions for unlocking this potential.

In light of point 2 above, we are currently focusing on micropower implementations of other parts of speech recognition systems, that are compatible with our analog pre-

processing system. We are currently evaluating prototype chips of these micropower implementation techniques, and will report early results in next quarter's status report.

3 Recent Publications

Lazzaro, J. and Wawrzynek, J., "Silicon Models for Auditory Scene Analysis," Neural Information Processing Systems (NIPS 95), December 1995.

Wawrzynek, J., Asanović, K., Kingsbury, B., Beck, J., Johnson, D., Morgan, N., "SPERT-II: A Vector Microprocessor System and its Application to Large Problems in Backpropagation Training," Neural Information Processing Systems (NIPS 95), December 1995.

Wawrzynek, J., Asanović, K., Kingsbury, B., Beck, J., Johnson, D., Morgan, N., "SPERT-II: A Vector Microprocessor System," to appear IEEE Computer, Spring 1996.